REPRESENTATION OF TIME-VARYING SHAPES IN THE LARGE DEFORMATION DIFFEOMORPHIC FRAMEWORK

Ali R. Khan and Mirza Faisal Beg

Simon Fraser University
School of Engineering Science
8888 University Drive, Burnaby, BC

ABSTRACT

Tracking and representation of shape change over time is of great interest in the field of computational anatomy. We propose a longitudinal growth model which estimates the diffeomorphic flow of a baseline image passing through a series of time-points that are the observed evolution of the template over time. We optimize the full space-time flow for the sequence of images, providing a linear space representation of the shape-change via a time-dependent velocity vector field, thus application of linear techniques becomes straightforward. We test our longitudinal growth model on both synthetic and real data-sets and demonstrate flexibility in time-point spacing, generation of average growth, and robust interpolation of missing time-points.

Index Terms— Longitudinal growth, Diffeomorphisms, Shape analysis, Computational anatomy

1. INTRODUCTION

Computational Anatomy involves developing mathematical algorithms for quantifying the extraordinary morphological variability of the human anatomy from 3D images acquired non-invasively using modalities such as MRI. These algorithms hold great promise for establishing a statistical understanding of 'normal' state, and have tremendous potential in early detection and treatment of diseases manifested by structural changes that do not fall within 'normal' limits observed in the population. Several cross-sectional databases now exist, and recently, databases containing longitudinal images acquired over time within the same subject have started to become available. These would enable the development of anatomical growth curves, that, akin to the growth curves so familiar for charting paediatric growth using simple measurements such as height and weight, would allow functional data curves representing evolution of anatomy over time to be developed and described statistically.

Quantifying shape change over time involves the analysis of a time-series of images that are sampled from the evolution of the anatomical state at discrete time-instants. To model anatomical evolution over time, Miller, Younes and Trouvé [1, 2] presented a dynamic growth model as the time-flow of a smooth velocity vector field carrying a known template image through a time series of images representing its evolution observed over continuous time; this however necessitates a dense sampling of anatomical evolution be available. Davis et al [3] have also proposed a kernel regression method to smoothly interpolate over time the cross-sectional data available at multiple time-instants.

In this paper, we revisit the longitudinal growth model framework, and estimate the time-flow of a template image through a sequence of discrete-time sampled images. This situation better represents the longitudinal databases being generated that contain within subject images over time that are sampled at variable intervals across subjects, and may not contain the same number of images for all the subjects. We derive the Euler-Lagrange equation for solving a variational problem that finds the optimal flow that smoothly fits the entire time-series. This is done in the large deformation diffeomorphic metric mapping (LDDMM) algorithm [4] framework, where the transformations for mapping a template image to a target image are modeled as the evolution of a smooth time-dependent velocity vector field. The integration of the estimated smooth flow, leads to the resulting transformations being diffeomorphic, which are smooth and invertible, with smooth inverse. The velocity space, V, in which the vector fields for constructing the flow are constrained, is a linear space and Hilbert, thus linear techniques and concepts can be applied directly to the velocity fields representing shape change over time.

2. METHOD

The diffeomorphic transformation connecting a time sequence of images is estimated via the basic variational problem in the space of smooth velocity vector fields V on domain Ω by minimizing the energy functional:

$$E(v) = \int_0^T \|v_t\|_V^2 \mathrm{d}t + \lambda \sum_{k=1}^{N} \|I_{t_0} \circ \phi_{t_k,0} - I_{t_k}\|_{L^2}^2. \quad (1)$$

Here, images are scalar functions on domain $\Omega \subseteq \mathbb{R}^d$; I_{t_0} denotes the template or baseline image, and I_{t_k} are N image time-points which make up the longitudinal time-sequence. The space V is Hilbert with inner-product $\langle f, g\rangle_V = \langle Lf, Lg\rangle_{L^2}$ defined using a differential operator L (see [5] for details). The evolution of the time-dependant velocity vector field flow, $v_t \in V$, is governed by the differential equation $\dot{\phi}_{0,t} = v_t(\phi_{0,t})$, thus for any $t \in [0,T]$ we get $\phi_{0,t} = id + \int_0^t v_\tau(\phi_{0,\tau})\, d\tau$. Hence, the matching transformation from baseline to time $t = t_k$ is found by integration of v_t from $t = 0$ to $t = t_k$, giving ϕ_{0,t_k}.

The minimizer of (1) is found via gradient descent of the variational gradient,

$$(\nabla_v E)_t = 2v_t - 2\sum_{k=1}^{N} K\big(1_{[0,t_k]}(t)\,|D\phi_{t,t_k}| \times (I_{t_0}\circ\phi_{t,0} - I_{t_k}\circ\phi_{t,t_k})\,\nabla(I_{t_0}\circ\phi_{t,0})\big), \quad (2)$$

where $K : L^2(\Omega,\mathbb{R}) \to V$ is a compact self-adjoint operator, which, by the construction of V, is uniquely defined by $\langle a, b\rangle_{L^2} = \langle Ka, b\rangle_V$. The indicator function is given by $1_{[a,b]}(t) = 1, t \in [a,b], = 0$ o.w.. Thus, the gradient in time interval $[t_i, t_{i+1}]$ is influenced by the images $I_{t_j} : j \geq i$, thus the flow between two image time-points is only influenced by the images forward or downstream in time.

2.1. Implementation

We implement our longitudinal growth model as an extension to the LDDMM framework, thus numerical implementation details can be found in [4]. The derived gradient is discretized and the update scheme, $v_t^{r+1} = v_t^r - \epsilon\nabla_v E_t$, is used to update the velocity vector fields at iteration $(r+1)$.

To deal with irregularly-spaced data in time, we discretize the velocity flow in relation to the time between scans; e.g. we can represent one month as one discrete time-step, thus images further apart in time will be represented by a proportional length velocity flow between them. The same strategy can also be used in the case of missing time-points.

3. EXPERIMENTS AND RESULTS

We performed several different synthetic experiments to show how our longitudinal growth model can handle irregularly spaced time-points, interpolate missing time-points, and construct an average growth sequence. We also demonstrate our model on real medical data, namely, the caudate nucleus in a Huntington's Disease patient.

The first experiment involves a dataset consisting of four sequences of a 3D sphere image ($64 \times 64 \times 64$) growing radially at different rates, with six time-points in each sequence (the baseline images are the same for each sequence). We apply the Longitudinal Growth Model for each sequence to obtain velocity vector field flows, then average these flows across the four different sequences to obtain the velocity vector field that corresponds to the average growth sequence, and forward integrate to get the average space-time maps. The velocity vector field was discretized using 10 time-steps between each image time-point. We also compute the longitudinal flows where a random $1/3$ of the time-points are removed from the computation. We generate a ground-truth average sequence using the average sphere growth rate, and compare that to the average obtained via the vector space averaging of velocity fields by comparing spherical volumes at each time-point using the error metric $(1/N)\sum_{k=1}^{N} |V(G_{t_k}) - V(\tilde{G}_{t_k})|/V(G_{t_k})$, where $V(G_{t_k})$ is the volume of the ground truth average sphere, and $V(\tilde{G}_{t_k})$ is the volumes of the computed average sphere at time t_k. The error for the average sequence computed using all data, and that from the set with incomplete data as compared to the theoretical averageis 8.95% and 9.81% respectively. Figure 1 shows the recovered average sequences for both these cases.

In the second experiment, we construct a synthetic dataset consisting of "typical" and "atypical" growth sequences of 2D (128×128) fish images. We use four different fish shapes and perform morphological dilations to generate a growth sequence of images with four time-points each. For half the fish shapes, we dilate the tail region at twice the rate to simulate atypical growth. Figure 2 shows these four synthetic growth sequences.

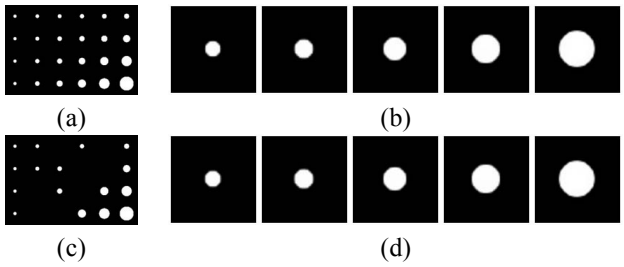

Fig. 1. (a) Four time-series synthetic 3D images showing a baseline sphere growing at different rates, (b) the average of these four sequences, (c) the original sequences with some intermediate samples missing, and (d) the average computed from these incomplete sequences.

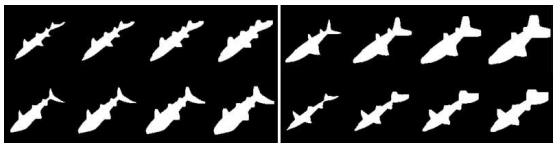

Fig. 2. Four synthetic time-series of images simulating fish growth; two typical growth sequences via uniform morphological dilation throughout the body of the template fish are shown on the left, and two atypical growth sequences with tail dilated at twice the rate as rest of the body is shown on the right.

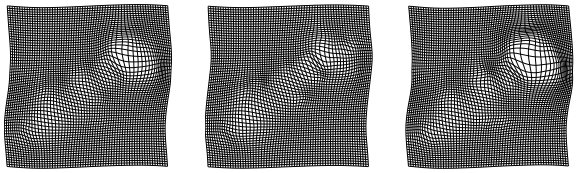

(a) All (b) Typical (c) Atypical

Fig. 3. The final transformation from the average longitudinal flow for the synthetic fish growth sequences using (a) all four time-series, (b) using only the typical group, and (c) using only the atypical group. The inflation in the tail region is captured in the average of the atypical sequences. Grids are shown down-sampled for clarity.

We apply the longitudinal growth model to each time-sequence of images to find the Hilbert-space representation for each growth sequence. We discretized the velocity field using 20 time-steps between each image time-point. Figure 4 shows the results of the longitudinal growth model applied to one of the normal fish growth sequences.

To further demonstrate robustness to missing time-points as shown in the growing sphere experiments, we apply the growth model to an atypical fish growth sequence with the third time-point missing; the resulting recovered image sequence is shown in Figure 5 . In this case, 40 velocity field time-steps were used between the second and fourth time-point to account for the missing image. As a simple comparison between typical and atypical fish populations, we apply averaging in the Hilbert-space vector fields and generate the average final maps for the typical, atypical and combined populations, as shown in Figure 3 .

For testing on real anatomical shape, we applied our longitudinal growth model to the caudate nucleus of a Huntington's Disease patient scanned on five separate occasions over a six year period. We used expert manual segmentations of the left and right caudate nucleus, smoothed with a Gaussian convolution filter (mask size: 5×5, $\sigma = 1$), and discretized the longitudinal time-flow proportional to the actual time-between scans as 20 time-steps per year. Segmentations of the left and right structures are combined into a single image volume ($54 \times 40 \times 34$) with image time-point rigidly registered to the baseline image. Figure 6 shows the evolution of the template image along the time-flow and the corresponding input data.

4. DISCUSSION AND CONCLUSION

Concluding, we have introduced a longitudinal growth mapping model which enables a Hilbert-space representation of time-dependent anatomical evolution in growth and disease. We have applied it to the mapping of volumetric images, although it can be generalized to other forms of data, such as landmarks [6], curves and surfaces as well. Our experiments show the model to be flexible with respect to temporal sampling of images and it is also able to interpolate missing time-points. The synthetic sphere growth experiment shows robust computation of average growth even in the absence of a significant amount of input data. The synthetic fish growth experiment simulated inter-subject variability in the presence of two distinct populations. We showed results on real medical data, the caudate nucleus in Huntington's Disease, for demonstrative purposes; we plan to apply our longitudinal growth model to clinical data in future longitudinal shape analysis studies involving statistical analysis of growth or atrophy extending similar methods that used two time-points [7].

5. REFERENCES

[1] M. I. Miller, A. Trouvé, and L. Younes, "On the metrics and euler-lagrange equations of computational anatomy," *Annual Review of Biomedical Engineering*, vol. 4, no. 1, pp. 375–405, 2002.

[2] M. I. Miller, "Computational anatomy: shape, growth, and atrophy comparison via diffeomorphisms," *NeuroImage*, vol. 23, no. Supplement 1, pp. S19–S33, 2004.

[3] B. C. Davis, P. T. Fletcher, E. Bullitt, and S. Joshi, "Population shape regression from random design data," in *IEEE International Conference on Computer Vision*, 2007.

[4] M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes, "Computing large deformation metric mappings via geodesic flows of diffeomorphisms," *International Journal of Computer Vision*, vol. 61, no. 2, pp. 139–157, Feb. 2005.

[5] A. Trouvé and L. Younes, "Metamorphoses through lie group action.," *Foundations of Computational Mathematics*, vol. 5, no. 2, pp. 173–198, 2005.

[6] M. F. Beg, M. I. Miller, A. Trouvé, and L. Younes, "The Euler-Lagrange equations for interpolating sequence of landmark datasets," *Sixth Annual International Conference on Medical Image Computing and Computer Assisted Intervention*, vol. 2879, pp. 918–925, November 2003.

[7] L. Wang, M.F. Beg, T. Ratnanather, C. Ceritoglu, L. Younes, J.C. Morris, J.G. Csernansky, and M.I. Miller, "Large deformation diffeomorphism and momentum based hippocampal shape discrimination in dementia of the alzheimer type," *Medical Imaging, IEEE Transactions on*, vol. 26, no. 4, pp. 462–470, 2007.

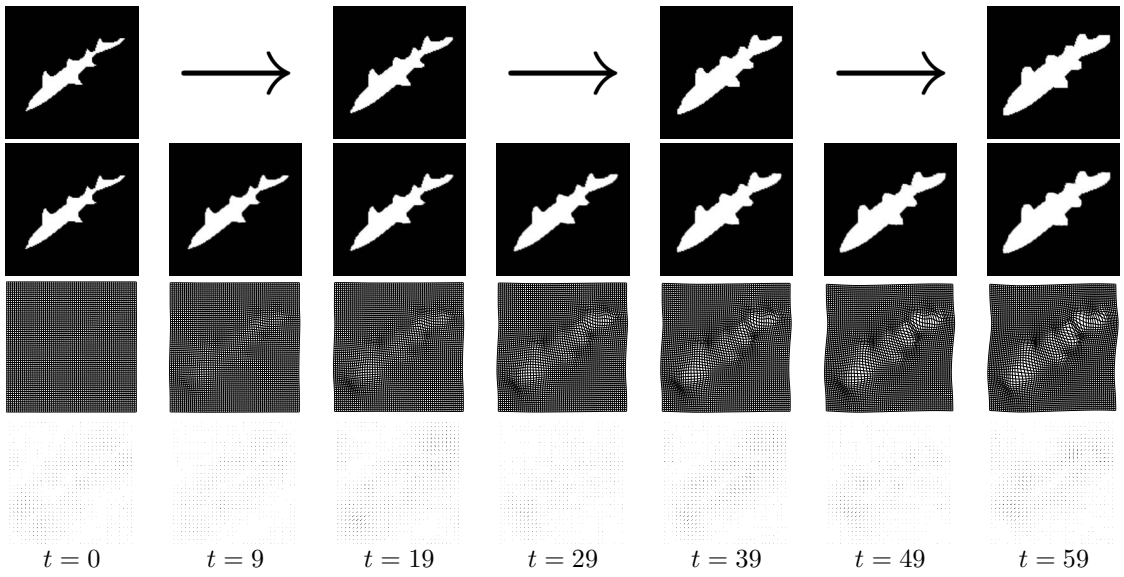

Fig. 4. Longitudinal growth model results for one of the typical fish growth sequences, showing one intermediate time-step computed between each input image time-point. From top to bottom we have the four time-points of the input synthetic data, the baseline image deformed along the generated longitudinal flow, the deformed grid corresponding to the transformation at the given time-step, and the the velocity vector fields. Grid and velocity vector fields are shown down-sampled for clarity.

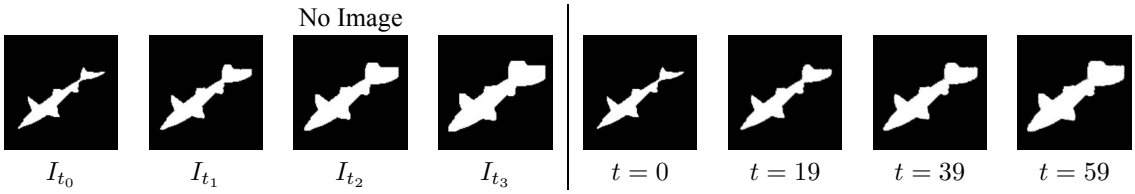

Fig. 5. Longitudinal growth model applied to abnormal fish growth with a missing input time-point (I_{t_2},$t = 39$). The input image sequence (including the missing time-point) is shown in the first row. The baseline image deformed along the longitudinal flow is shown in the second row; the interpolated image for the missing time-point closely matches the actual image.

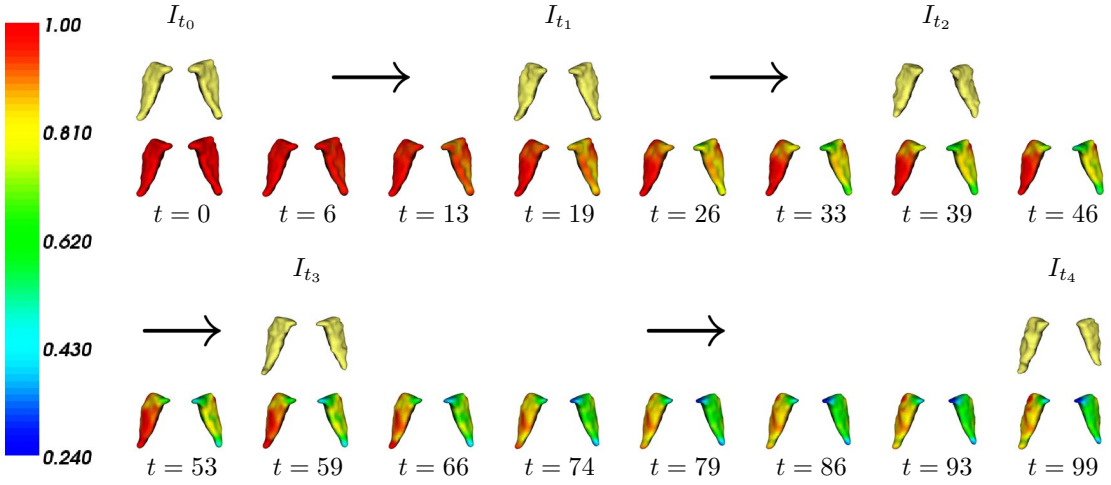

Fig. 6. Longitudinal growth model results for the Huntington's Disease caudate nucleus where the subject was scanned five times over the course of six 6 years; the time-flow is discretized to satisfy 20 time-steps per year. Input shapes are shown along with the baseline shape as deformed along the flow, colored with $|D\phi_{t,0}|$, with values below 1 indicating localized loss of tissue.